

ESD ACCESSION LIST

XRRI Call No. 82138

Copy No. 1 of 2 cys.

Technical Note

1975-8

Preliminary Investigation
of Digital Speech Processor
Hardware Implementations

P. E. Blankenship

5 February 1975

Prepared for the Defense Communications Agency
under Electronic Systems Division Contract F19628-73-C-0002 by

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS



Approved for public release; distribution unlimited.

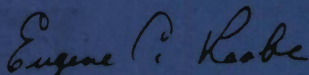
ADA007062

The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology, for the Military Satellite Office of the Defense Communications Agency under Air Force Contract F19628-73-C-0002.

This report may be reproduced to satisfy needs of U.S. Government agencies.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER



Eugene C. Raabe, Lt. Col., USAF
Chief, ESD Lincoln Laboratory Project Office

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

PRELIMINARY INVESTIGATION OF DIGITAL SPEECH
PROCESSOR HARDWARE IMPLEMENTATIONS

P. E. BLANKENSHIP

Group 24

TECHNICAL NOTE 1975-8

5 FEBRUARY 1975

Approved for public release; distribution unlimited.

LEXINGTON

MASSACHUSETTS

ABSTRACT

The present Lincoln Laboratory Digital Voice Terminal (DVT) is analyzed with the intent of improving form factor and cost figures. It is found that, with standard integrated circuits, improvements are possible only if a substantial performance penalty is paid and that the present configuration represents a creditable cost/performance trade-off. Utilization of custom LSI is rejected at this time as being too expensive an approach given the modest level of production expected. A hybrid packaging technique is seen to improve form factor at a much lower developmental cost/part than custom LSI and may be of interest for low level production. Semi-programmable architectures, based on commercially available bipolar LSI μ -processor chip sets, seem to afford a very promising near-term solution to the low cost, mass producible narrow-band voice terminal design problem. Performance levels are sufficient for the projected computational loads though the overall speed and flexibility of a DVT-like structure are largely sacrificed.

I. Introduction

Given the continued interest in digital compressed speech expressed by various governmental agencies, it seems worthwhile to expend some effort on the problem of searching out innovative, efficient processor structures which take advantage of present-day technology evolutionary trends. The quest should focus upon candidate designs which appear promising from the following viewpoints:

1. low unit cost
2. amenable to high volume production
3. high reliability
4. compact form factor
5. flexible/versatile architecture

Tradeoffs in emphasis amongst the (potentially conflicting) desired objectives yield designs which can be roughly classified into three fundamental categories:

1. Special purpose: this approach typically embodies the most efficient, compact, and inexpensive approach to implementing a particular choice of algorithm. The price that is paid, of course, is the relative inflexibility of the end product.

2. General purpose: This class of processor, since it incorporates what basically amounts to a computer, is by virtue of its wholly programmable nature the ultimate in terms of flexibility. However, for a specific algorithm choice, inevitable inefficiencies imply a tougher overall performance requirement with all of the attendant problems indicative of high speed technology system design. Stated simply, for a given problem the design is bigger and more costly than it probably need be.

3. Hybrid: in the expansive middle ground lying between the aforementioned extremes, exists a necessarily broad spectrum of designs which attempts to marry the best aspects of both worlds. Such hybrid designs are partly special purpose and partly programmable. For example, a functional building block common to many processing schemes (like correlation), but which is particularly taxing computationally, might be built as a special purpose subsystem. But complicated specialized tasks, such as reflection coefficient extraction in an LPC vocoder analysis, might best be implemented in a limited programmable section.

It is our contention that a high premium should be placed on the more flexible design alternatives for active research applications areas such as speech processing. Given the many systems already in existence (APC, LPC, Channel, VELP, etc.) and the many more which will no doubt evolve, a fully flexible research vehicle seems essential. The first part of this report focuses upon a recently developed entry into the programmable processor category: The Lincoln Laboratory Digital Voice Terminal. The intent is to suggest possible methods of reducing the cost and improving the form factor of the current design. Upon careful scrutiny the design is found to be dominated in terms of cost, integrated circuit count, and performance by its extremely fast memory complement. It is shown that little can be done to improve the design if constrained to maintaining the current performance levels with standard integrated circuitry. It is further shown that a slower version which utilizes less expensive, more dense memory chips can be had at a 30% decrease in circuit count at a 50% speed penalty. Switching to a lower speed technology is found to afford a similar package count reduction at a 100% penalty in speed, but the overall cost per unit is halved over the current design.

High performance technology custom large scale integration (LSI) is evaluated as an alternative to off-the-shelf parts. It is seen that this approach, which does not address itself to the memory area since that is considered a specialty, can be expected to impact little with respect to IC count on the current memory dominant design. Custom LSI is also found to be expensive in terms of developmental costs per unique part type. For low volume production, such expenses cannot be justified.

A hybrid packaging scheme, wherein several dice of standard, off-the-shelf design share a common substrate, is suggested as a reasonable compromise approach. The developmental costs per part are about 1.5 orders of magnitude cheaper than LSI and the memory density issue can also be accommodated. Form factor and reliability improvements similar to those of genuine LSI can be expected though the raw cost of IC dice as supplied by the vendors does not drop appreciably over that of standard packaged units.

The second part of the report concerns itself with the application of newly available bipolar LSI microprocessor chip sets to the problem of speech processor design. It is shown that the devices are by themselves far too slow to compete with DVT-like performance and that programmable parallel processing architectures based upon them do not yield satisfactory results in terms of utility, cost, form factor improvement, or performance. Hybrid or quasi-programmable processor structures are suggested as likely application candidates for the microprocessors. One such structure, specialized to the task of LPC processing, is described. Initial estimates of integrated circuit count and attendant costs are indicated.

II. General Purpose Processor Case Study: The DVT

Lincoln Laboratory has designed and constructed a speech processor of the general purpose class called the "Digital Voice Terminal" (DVT). The heart of the system is a custom designed 16-bit, 2s complement, fixed point programmable processor comprised of about 470 two-nsec emitter coupled logic (ECL), medium scale integrated (MSI) circuits. The basic execution cycle is 55-nsec for all operations (excepting multiplication which requires 220-nsec) putting the processor in the 18 mega-instructions per second category. The instruction set is fully flexible, containing effectively 128 operations, and is alterable through a micro-code ROM such that it can be tailored and optimized to specific tasks.

The memory complement consists of a 25-nsec access 512 x 16-bit data memory (M_D) and a separate 1024 x 16-bit program memory (M_P) containing executable code exclusively. Both are realized with high performance ECL bipolar technology. An overall block diagram of the programmable processor is shown in Figure 1.

To specialize it to the task of speech processing, the programmable processor is connected through a versatile in-out structure to a collection of physically integral peripheral devices (Figure 2). One data path is connected to an A/D-D/A converter set which, along with its associated sampling and filtering hardware, drives the user handset. A second path is devoted to a serial-to-parallel/parallel-to-serial converter set which mediates modem traffic flow over phone lines (or wireless transmission mechanisms) to other speech processors. A third path, optimized in speed, connects to an auxiliary, fast, 2048 x 16-bit bipolar random access memory which serves to enhance the programmable processor's internal data storage capacity. Yet another path is connected to a non-volatile program memory image which can be loaded into M_P automatically on power-up if the DVT is operating in a stand-alone rather than a laboratory environment. The DVT can also

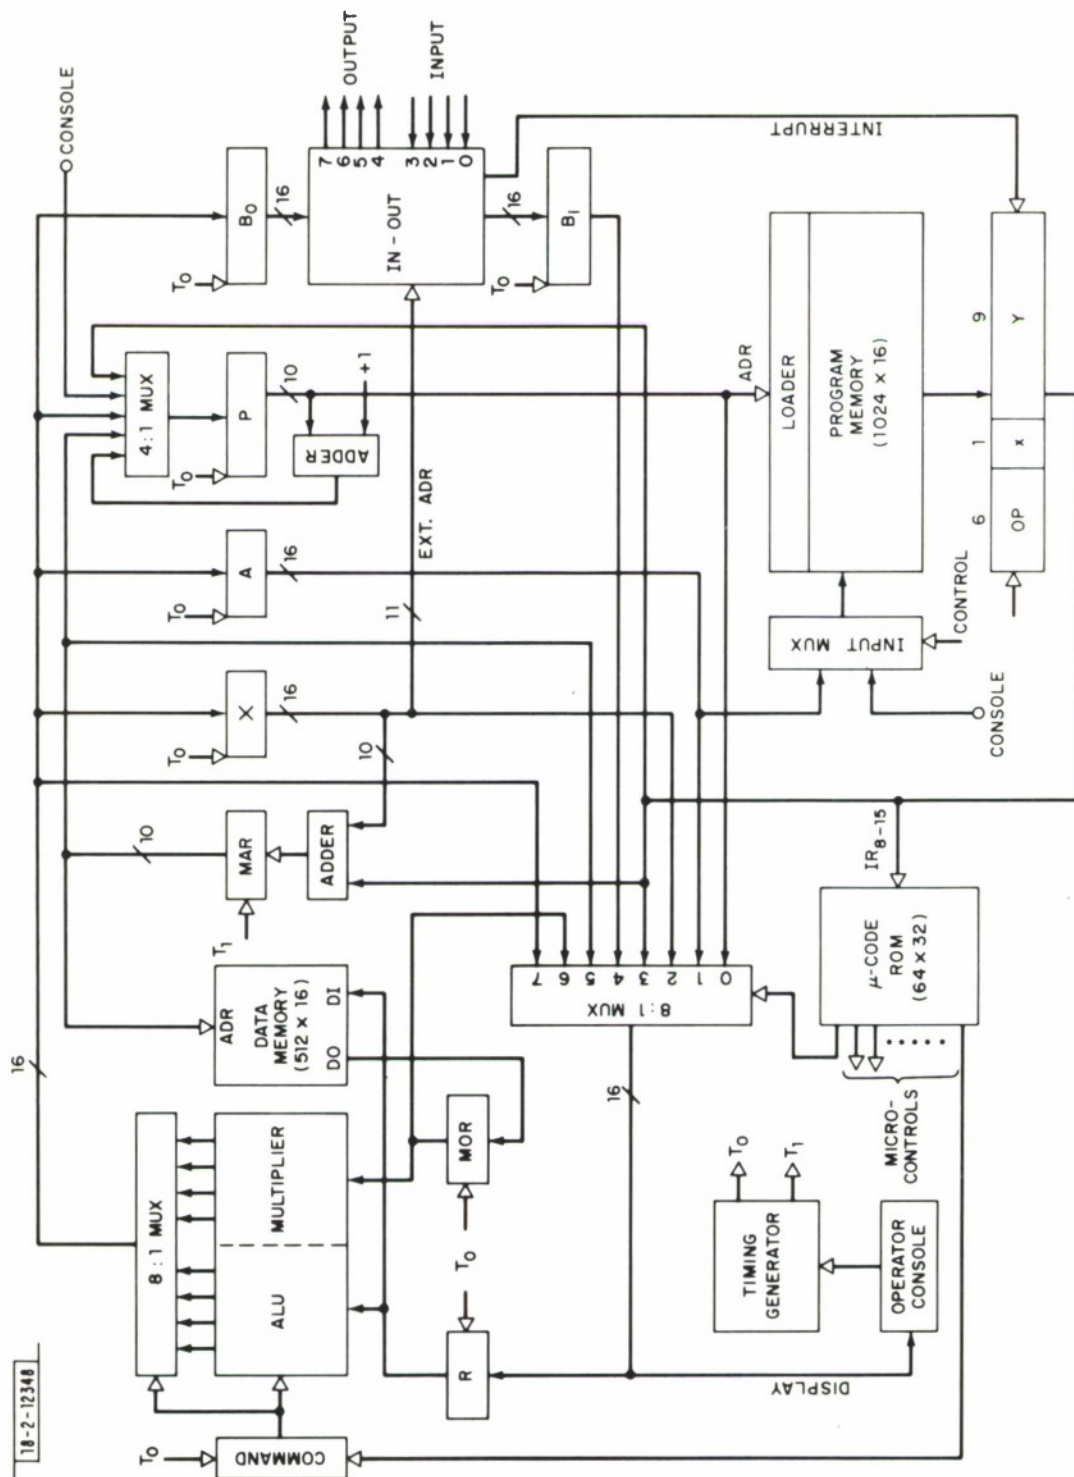


Fig. 1. DVT Programmable Processor

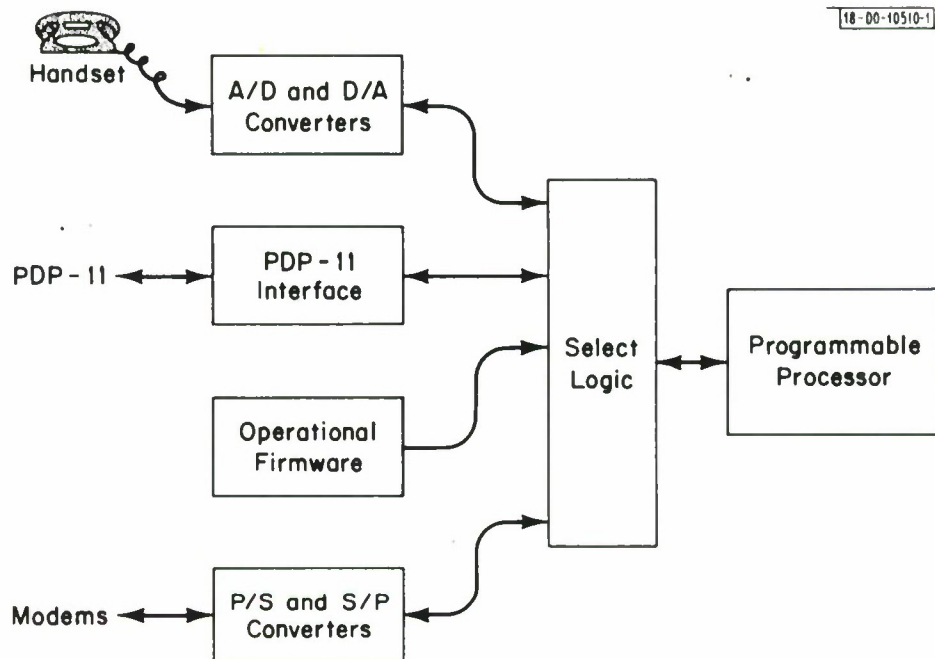
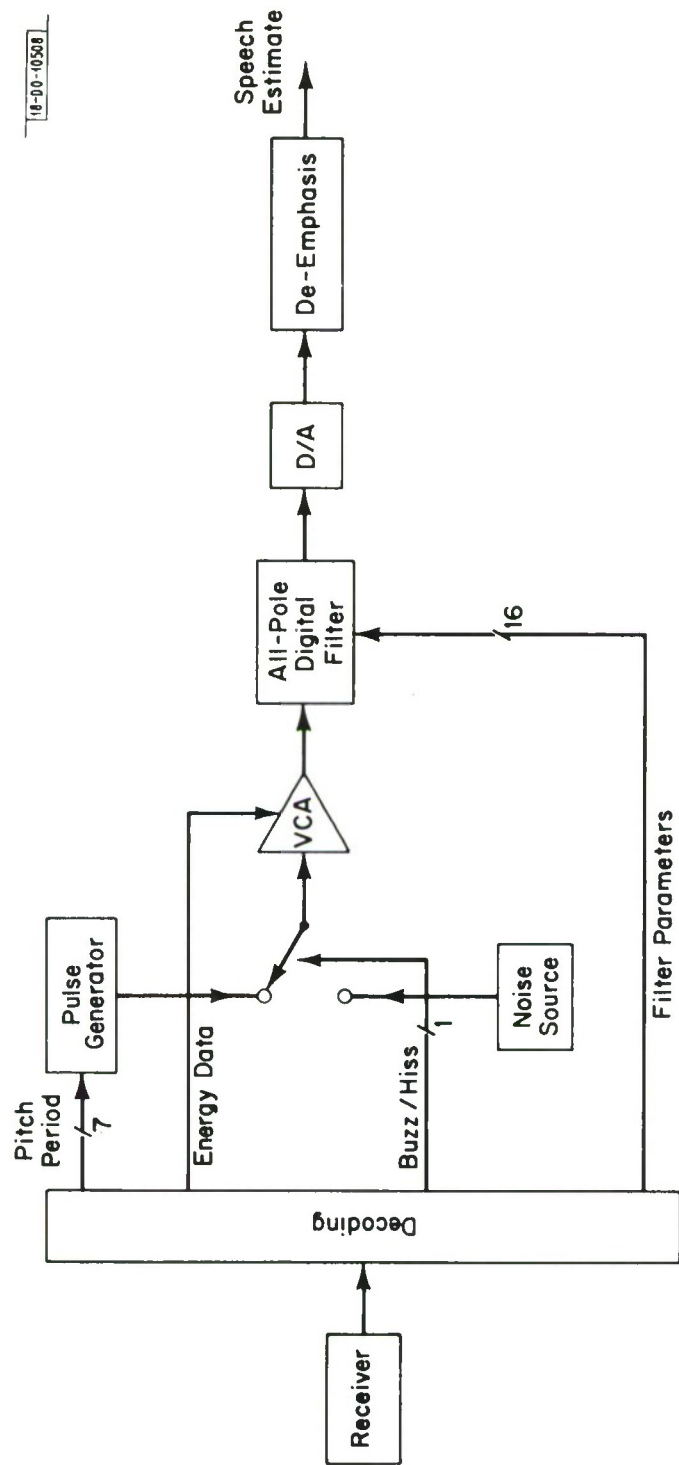


Fig. 2. DVT Input/Output Structure

be operated in conjunction with a host major data processor, if desired, and an inter-computer I/O channel is provided for this purpose. The I/O hardware features a minimum latency, vector interrupt capability which insures rapid response and maximum real-time programming ease. About 125 saturating logic circuits (TTL) are required for the peripherals.

To assess the DVT's performance in a practical situation, the essential software components of a 12th-order Markel LPC¹ vocoder system have been coded as a benchmark. The synthesis scheme, shown schematically in Figure 3, centers upon an all-pole time-varying filter as a model of the human vocal tract. The filter is excited by either a white noise source, or a pulse generator controlled by the transmitted pitch period estimate, depending on whether a given frame is voiced or not. The more complex problem of analysis is shown schematically in Figure 4. Parameters characterizing the vocal tract model for a given speech frame are extracted via an autocorrelation followed by a Levinson recursion². Asynchronous pitch estimation is conducted in parallel using the Gold-Rabiner method³. The 12 filter parameters, voice energy level estimate, buzz/hiss decision, and pitch period estimate are finally encoded and packed for transmission.

Computation time estimates for the various requisite processing tasks are listed in Table 1. Each task is categorized as to whether it belongs to analysis or synthesis, and whether it must be performed once per speech sample, or once per frame. The table was compiled assuming a sampling rate of 6.6 KHz, and 22.5-msec speech frames overlapped by 33% which is equivalent to an intersample period of 150 μ sec and an effective frame rate of 67 Hz. The autocorrelation time assumes double precision arithmetic and that 2 correlation updates are performed on each sample arrival. Based on this information, it seems that the DVT is capable of exceeding real time by about 100% for this LPC implementation.



16-00-10508

Fig. 3. LPC-12 Synthesis

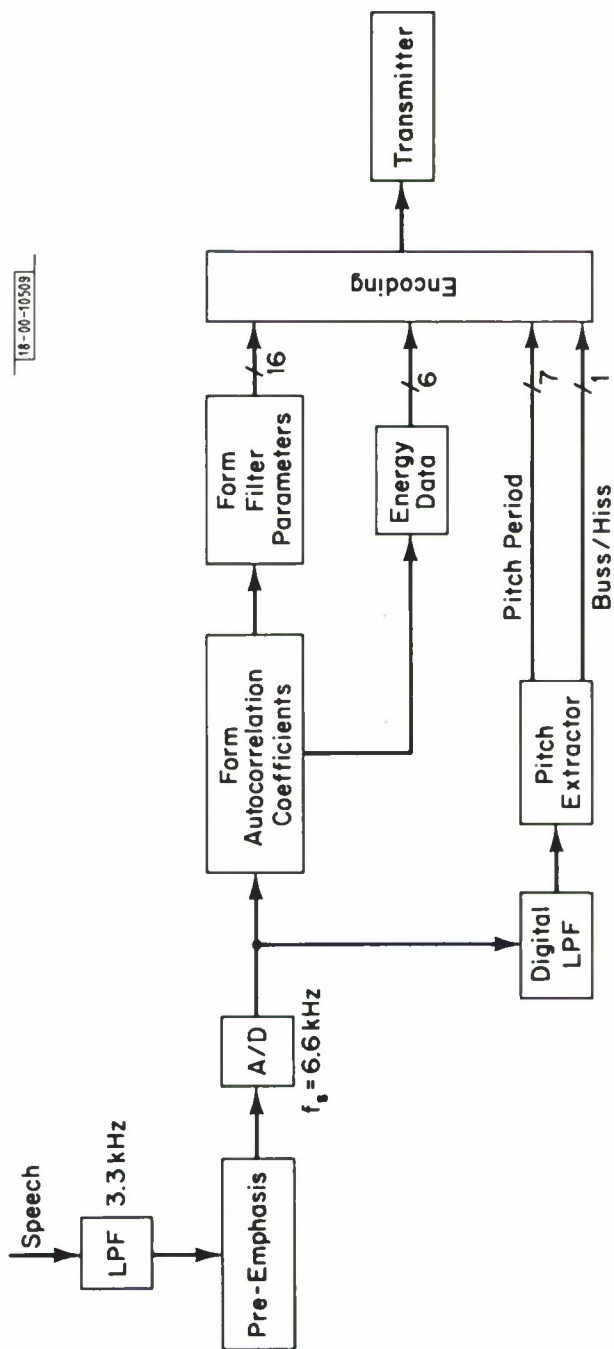


Fig. 4. LPC-12 Analysis

TABLE 1
MARKEL LPC-12 REAL-TIME PERFORMANCE

	COMPUTATION	T _{COMP} (MICRO-SECONDS)	
		PER SAMPLE	PER FRAME
ANALYSIS	CORRELATION AND WINDOW	20	--
	FILTER PARAMETER EXTRACTION	--	262
	PITCH DETERMINATION AND BUZZ/HISS DECISION	35	275
	PARAMETER ENCODING	--	88
SYNTHESIS	PARAMETER DECODING	--	13
	BUZZ/HISS GENERATION	1.6	--
	FILTERING FUNCTION	11.1	--
	TOTALS	67.7	638
	$T_{\text{COMP.}} / T_{\text{AVAIL.}} = \frac{67.7 = \frac{638}{100}}{150} = .49$		

In order to assess what might be done to improve package count and cost, it is interesting to see how the DVT's nominal 470 ECL IC allotment and \$13,000 outside purchase budget was spent. Table 2 shows a listing of the programmable processor's major subassemblies and the ECL circuit count associated with each. A striking observation is that something over a third of the circuits were used up in the 2 internal memories. In terms of dollars, these 2 items comprise about 2/3 of the overall circuit cost for the programmable processor. Table 3 summarizes these facts.

Table 4 enumerates in some detail the recurrent outside purchase (O.P.) charges sustained by Lincoln Laboratory related to the production of a single DVT unit. These figures do not reflect overhead associated with design, fabrication, and debug of each unit. Total integrated circuit costs comprise about 42% of the total with the ECL accounting for a full 28%. If the ECL memory alone is examined, it is seen that these circuits comprise nearly 20% of the total. It is also interesting to note that wire-wrap charges plus the requisite circuit panels, wire, terminations, and decoupling capacitors amount to 20% of the total -- as much as the entire ECL circuit cost! These observations reflect the cost penalty associated with a high performance wire-wrap system. If a commercial vendor were to implement the current design with a very modest production level projection (\approx 100-1000 units), he would attempt to minimize his costs primarily by:

1. obtaining quantity discounts on digital and analog semiconductor components
2. using multilayer PC boards (\sim 4 signal layers) instead of wire-wrap.

Estimates indicate that, for a commercial DVT, the \$13,100 figure (Table 4) would drop to about \$8,200 given the above considerations.

Table 5 suggests some minor design revisions which essentially retain current processor performance while permitting some small circuit count reductions. It is possible to shave a few circuits off the arithmetic

TABLE 2
DVT ECL PACKAGE COUNT BREAKDOWN

<u>SUBSECTION</u>	<u>16 PIN</u>	<u>24 PIN</u>
P REGISTER	28	0
INSTRUCTION REGISTER	22	0
CONTROL DECODING	14	0
INPUT/OUTPUT	45	0
CLOCK GEN. & CONSOLE CONTROL	30	0
ALU	39	4
16 x 16 MULTIPLIER	44	8
M _D ADDRESS CONTROL	29	0
R REGISTER GATING	23	0
1024 x 16 PROGRAM MEMORY	86	0
512 x 16 DATA MEMORY	83	0
MISCELLANEOUS	21	0
	<hr/>	<hr/>
	464	12

TABLE 3
DVT ECL MEMORY COST BREAKDOWN

ITEM	COUNT	% COUNT	COST	% COST
M _P	135	18	\$1350	36
M _D	135	18	\$1000	27
OTHER	300	64	\$1250	37
TOTAL	470	--	\$3700	--

TABLE 4
DVT SUBASSEMBLY COST
BREAKDOWN

ITEM	COST	% COST
ECL CIRCUITS	\$ 3700	28
TTL CIRCUITS	1800	14
ANALOG DEVICES	700	5
POWER SUPPLIES	1000	8
WIRE-WRAP PANELS	2000	16
WIRE-WRAP CHARGES	600	5
RESISTORS/CAPACITORS/WIRE	950	8
CONNECTORS	700	5
ENCLOSURES	650	5
MISCELLANEOUS	1000	8
TOTAL RECURRENT O.P. COSTS PER UNIT	\$13,100	

section by removing some shift multiplexing and using a new multiplier chip which is due from Motorola in first half 1975. Some control revisions, such as register clock gating in lieu of recirculation, also save a bit. But in all, a reduction of only about 50 circuits seems possible.

Clearly, in order to realize any appreciable package count and cost improvements it is necessary to attack the memory dominance issue. Memory densities increase and cost/bit decreases as performance requirements are relaxed. Table 6 suggests some design revisions which take advantage of cheaper memory at a penalty in overall processor performance. Item #2 implies a resident non-volatile program memory (ROM) and precludes operating the DVT in anything but a stand-alone mode. Items 3 and 4 retain the present random access memory structures in M_P and M_D but use slower memory devices. As it happens, a minimal performance penalty is suffered in changing the processor's timing philosophy from a triple overlapped to a double overlapped arrangement while saving some additional control circuits. This can be seen by comparing the cycle times of items 3 and 4. The net result is that essentially the same processor structure can be retained while eliminating about 1/3 of the integrated circuits at a performance penalty of 51%. Since the LPC-12 benchmark program appears to run at half real time, such a performance degradation would appear easily tolerable for this application at least. In terms of money, the ECL components cost would be reduced to about \$2900: an improvement of 20%.

If a factor of 2 in performance degradation can be withstood, it seems reasonable to consider a technology shift to a saturating logic family such as the standard 54/7400 series TTL MSI. There is ample motivation for doing this since parts and fabrication costs can be drastically reduced. An improvement in form factor can also be expected because much more compact power supplies could be employed. (About $\frac{1}{2}$ of the 1.25 ft^3 volume occupied by the DVT is power supply.) Calculations indicate that a TTL design

TABLE 5
REFINEMENTS OF CURRENT DESIGN

DESIGN REVISIONS	IC's SAVED
1. Use MC 10183 in DVT multiplier.	17
2. Gate register clocks instead of recirculate.	15
3. Use Hex D (10176) flip-flops & Hex (10195) inverters in clock gen.	9
4. 4 ALU output options instead of 8.	9
TOTAL	50

TABLE 6
ALTERNATE DVT DESIGNS

DESIGN ALTERNATIVES	CYCLE TIME	IC's SAVED
1. Triple overlap with RAM M_p .	55.0	50
2. Triple overlap with 1K x 40-bit ROM as M_p .	55.0	96
3. Triple overlap with slow M_p (F10415) & M_D (F10410).	81.3	130
4. Double overlap with slow M_p & M_D .	83.0	152

corresponding to item 3 of Table 5 would exhibit the same package count as the ECL version at an integrated circuit cost savings of about 50%. This is primarily due to the relatively inexpensive TTL memory chips. System design cost savings are realized also in such areas as circuit panels, terminations, power supplies, power supply decoupling, and metal work. Rough calculations indicate an overall fabrication cost savings of about 50% can reasonably be expected. However, a 110-nsec cycle time design is not possible with standard TTL. Upwards of 130-nsec is a more reasonable estimate. It would be necessary to make use of a limited number of judiciously selected high speed TTL circuits (Schottky series) to attain a 110-nsec cycle time goal. This complicates the system design and increases the power budget somewhat thereby compromising expected savings in these areas.

With the advent of several viable bipolar large scale integration (LSI) technologies⁴, it is informative to consider their implications with respect to the current DVT design. LSI implies in present day terms 500 to 10,000 devices per chip. Some rather obvious advantages of this philosophy are:

1. minimum system size, weight, and power dissipation
2. fewest number of chips per design
3. high reliability due to decreased number of IC interconnects
4. improved maintainability
5. improved performance potential due to minimized interconnect lengths
6. for high volume production, recurring fabrication costs per unit are minimized

The disadvantages are simply the high developmental cost and relatively long design cycle time per unique part type. Expenditures on the order of \$50,000 to \$100,000 per chip design and turnaround times on the order of 9 to 12 months are not unusual.

To specify a custom family of LSI chips for the DVT with a minimum number of unique part types, the existing design must be partitioned in an

optimum manner. In order to do this effectively, it is desirable that the design exhibit a regular or iterative topology. If it turns out that this is not the case, it is necessary to define very complicated, cumbersome ships to keep the number of part types under control. Such chips characteristically fall into what is termed "very large scale integration" (VLSI) technology which implies more than 10,000 devices per chip. Such complexity is beyond the present day capabilities of ECL technology, but some work of this type has been done with the much lower performance emitter follower logic (EFL)⁵. However, because of the decreased device performance of this technology, it does not seem possible to construct a DVT-like processor that can even meet real time requirements let alone match its performance.

Upon examination of the current design, it is seen that only the arithmetic and register file sections exhibit any apparent regularities. The very fertile area of memory is explicitly excluded since no custom LSI house that we know of is doing work in this area. A four-bit slice through the register file was considered but pin out requirements imply a large header (at least 28 pins). Since only 10% of the total package count is tied up in this subsystem, LSI would have negligible overall effect here anyway. The adder/subtractor, using efficient MSI chips, is highly integrated already. The multiplier, however, could benefit from LSI both in local package count and performance potential though the overall system form factor is not drastically improved. A 4 x 4-bit, 2s complement multiplier chip currently under development by Lincoln Laboratory is shown in Figure 5. It is realized with a higher-than-standard performance ECL technology and can be packaged in a 28-pin header. Incorporated into the current DVT design, it would save 25 16-pin packs and replace 8 24-pin packs with 4 of the 28-pin class. An attendant 25% improvement in multiplier performance can also be expected.

Fig. 5. 4 x 4-Bit 2s Complement Multiplier

A less costly approach to form factor improvement which encompasses several of the benefits afforded by LSI and yet can be applied to the memory issue is termed "hybrid packaging". With this method, standard die as supplied by the manufacturer are bonded to a common substrate. Chip interconnects are effected by wire bonds to single layer substrate metalization. Performance, reliability, and even dissipation (due to reduced load capacitance seen by on-chip drivers) can be improved, not to mention a repairability feature. Developmental costs are on the order of a few thousand dollars per part type and design cycle times are on the order of several weeks.

As a typical example, a 128 x 8-bit memory package, currently under development by at least one vendor, is shown in Figure 6. The design is based on a fast ECL 128 x 1-chip which accesses typically in 11-nsec. This particular configuration, containing 11 die and dissipating about 5W would substitute 8 28-pin packs for the 83-odd 16-pin packs which currently constitute M_D . A similar strategy could be formulated using the ECL 256 x 1-memory chip yielding similar savings in the M_P design. Raw integrated circuit component costs do not improve with this technique, however, since manufacturers charge virtually the same for dice as for a packaged unit (based on charges for a molded plastic commercial header). But a real estate improvement of about 5:1 is realized in the M_D and M_P subsystems. Power dissipation density is certainly increased but forced air coupled with miniature heat sinks still is a viable cooling approach.

From the foregoing discussion, the following conclusions are drawn with regard to the current DVT architecture:

1. Given the degree of performance desired, the constraints of a standard package design, and a tight schedule, the choice of ECL technology in a wire-wrap environment was essential and the final package count, size, weight, and cost not unreasonable.

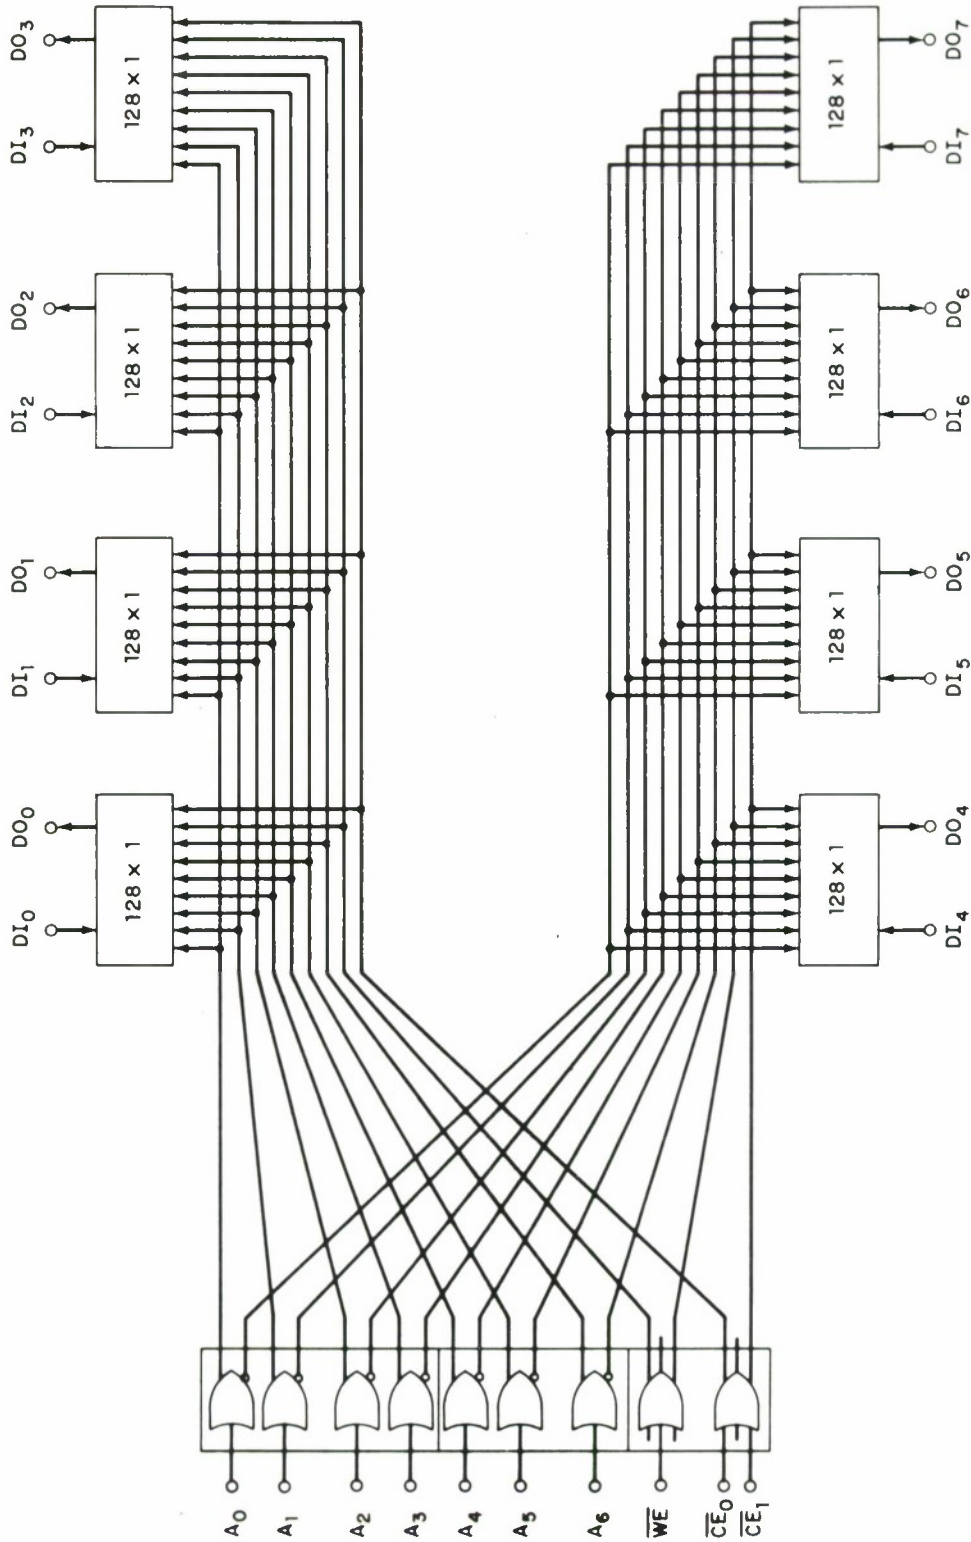


Fig. 6. 128 x 8-Bit ECL Memory Hybrid Pack

2. No significant improvement in package count, performance, and form factor is possible with currently available standard ECL integrated circuits.

3. Significant package count reductions are only possible with a marked overall performance degradation. This is primarily due to constraints imposed by the bipolar memory dominance of the design in both cost and integrated circuits. Use of higher density memories which exhibit a lower cost/bit and concomitant performance degradation impacts heavily in both these areas.

4. Switching to saturating logic technologies for low performance options would cut overall costs in half and still yield a processor which is a factor of 4 or 5 faster than those commercially available. However, it is not clear to us that processors of the DVT architectural ilk in this performance class are of high prospective utility as speech research tools given the uncertainty in complexity and computational onus of future processing schemes.

5. Due to the nature of the DVT architecture and the performance level demanded, it does not seem possible to define a small number of unique LSI parts, with complexities not beyond the realm of ECL technology, which would have more than a token impact on system IC count. Given the high developmental costs/part type and the relatively low level of DVT production expected, custom LSI should probably be rejected as economically unfeasible.

6. The hybrid packaging approach does seem to exhibit a potential for overall system form factor improvement, even in the memory area. Though apparently no dollars are saved in IC die procurement, the developmental cost/part type are at least an order of magnitude more palatable than the LSI approach. However, the recurrent fabrication costs per piece may prove to be prohibitive since this is a very laborious process. Therefore the hybrid

technique should be investigated further, but cautiously, for memory dominant, low production volume designs such as the current DVT.

III. Quasi-Programmable Processors Using Bipolar Microprocessor Elements

Within the last year, 2 relatively low cost bipolar microprocessor chip sets have become available as standard offerings and it appears that at least 2 additional manufacturers will be entering the market place in the near future. These circuits legitimately qualify in complexity as being of the LSI class and are realized with a form of Schottky TTL technology. Applications areas which can withstand the performance limitations inherent in such devices can avail themselves of the following obvious advantages:

1. As was shown earlier, a TTL system design is on the whole cheaper and less complex than a high performance ECL system.
2. LSI componentry affords many advantages yet the exorbitant cost of devising custom parts for a particular design is avoided.

The LSI units described here impact greatly on what would normally be considered the arithmetic and control portions of a standard mini-computer architecture. They rely heavily upon recent advances in bipolar read-only-memory (ROM) manufacturing technology and do not address the issue of random access memory at all.

These chip sets are designed to be used in the context of a micro-programmed architecture⁶, a typical form of which is shown in Figure 7. The advantage of the micro-programming concept is that the character of the processor (i.e., the effective instruction set) is defined by the contents of a ROM. Therefore a single general logic structure can, if fast enough, be made to look like (or emulate) any existing computer design from the user's viewpoint. The canonic architecture consists of a central processing element (CPE), a control, an input/output section, and a main random access store for both code and data. The cleverness of this arrangement is embodied in the control, which is comprised of sequencing logic and the characteristic ROM. Each complex computer instruction is decomposed into a sequence of elemental steps (μ -instruction) which are contained in the ROM. The micro-

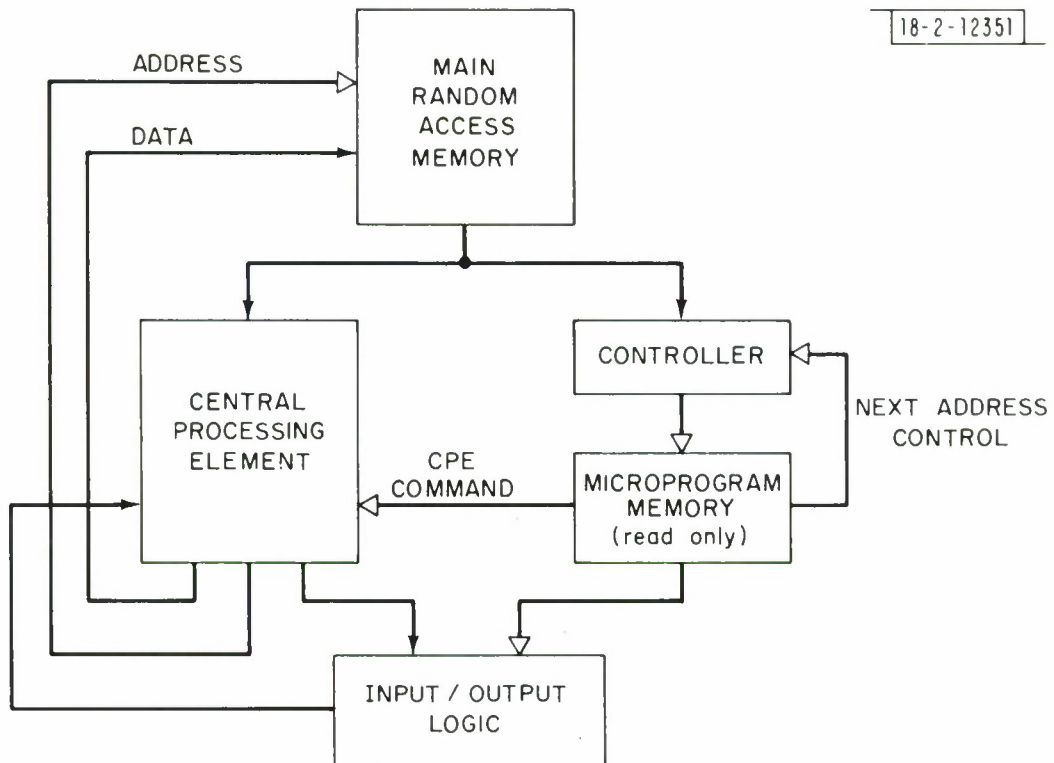


Fig. 7. Typical Micro-Processor Architecture

program controller sees to it that each micro instruction is executed properly in sequence and that new complex (or "macro") instructions are fetched from the main store at appropriate times. In actuality the ROM is the key element in the design since it replaces much of the bothersome random logic characteristic of computer controls.

Block diagrams depicting the essentials of the two existing CPE elements are shown in Figures 8 and 9. The unit of Figure 8 consists of a 2-bit slice through an arithmetic/logic unit (ALU), an 11-deep scratch pad register file, an accumulator, and an auxiliary buffer register.⁷ Attendant decoding and selection logic is also provided locally on the chip. In a 16-bit context this element is capable of 120-nsec clocking epochs for elemental μ -instructions such as an addition involving the accumulator and the scratch pad file. However, to perform a typical macro-instruction, several elemental cycles may be required. A typical sequence for an addition between a scratch pad register and a location in main memory might proceed as follows:

1. Compute effective memory address and store in address register.
2. Load memory into accumulator (AC).
3. Add scratch pad register to AC and store.
4. Increment program counter and load address register.
5. Load next macro-instruction into instruction register from main memory.

Thus, 5 elemental epochs are necessary to perform one macro-instruction and fetch the next, a total time of $5 \times 120 = 600$ -nsec. This is about a factor of 10 slower than the DVT. More complex operations such as multiplication can, unless special hardware is added, take up to 20 times longer than the DVT. Given that the architecture of this CPE is not terribly dissimilar to that of the DVT, it seems apparent that even 2 such micro-processors operating in parallel (one for analysis, one for synthesis) cannot vaguely approach the performance levels of the DVT for LPC.

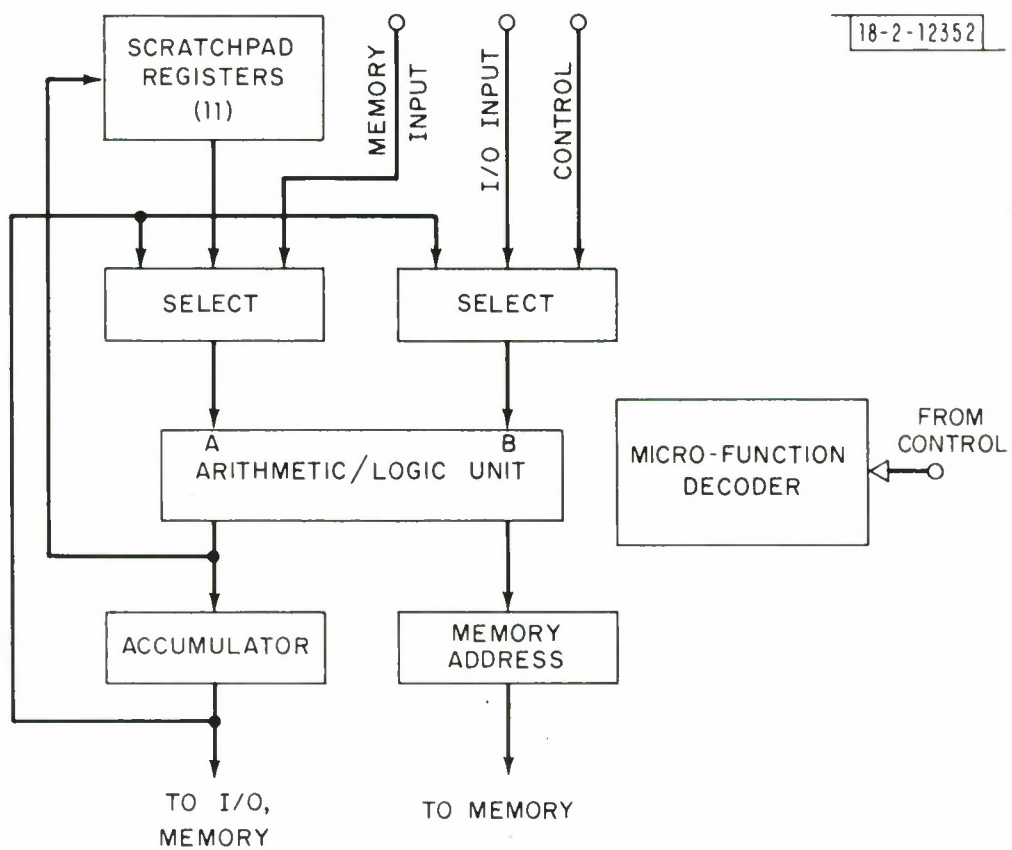


Fig. 8. Single Address Register File CPE Chip

18-2-12353




Fig. 9. 2-Address Register File CPE Chip

For completeness, a second type of CPE is shown in Figure 9. It consists of a 4-bit slice through an ALU, a 16-deep 2-address register file, and an auxiliary register. Local decoding and selection logic is included. In a 16-bit context this unit is capable of a 200-nsec μ -cycle epoch. Though apparently slower than the other CPE element, this unit features architectural advantages which could, in some applications, offset its relative sluggishness. The 2-address register file could reduce main memory accesses thereby speeding up overall execution times. To test this thesis, the Levinson recursion portion of the DVT's LPC analyzer was coded on a paper-design processor based on this CPE. The design of interest employed much auxiliary external logic to reduce the number of μ -cycles per macro-op to the bare minimum (namely 1). Even so, the execution time turned out to be no better than the ratio of its clocking epoch to the DVT's. Hence, it was concluded that the 2-address cache memory does not afford any obvious advantages in this case and the overall performance of this CPE could be expected to be even worse than that of the other for a full LPC. Another disadvantage of this element is that it is the only member of its chip set. The set which complements the 1-address CPE contains a μ -controller, look-ahead carry block and priority interrupt in/out control.

Returning for a moment to the notion of paralleling microprocessors to achieve equivalent performance to the DVT's, it is interesting to pose the question: Where is the point of diminishing returns? This query can be dealt with summarily by considering the case of 4, parallel 1-address processors sharing, perhaps, a common main store. The following conclusions can be drawn from studying such an arrangement:

1. Though as general as the DVT, this is a far more difficult structure to coordinate and program.
2. In terms of performance this arrangement is still, on the average, $11/4 = 2.75$ slower than the DVT.

3. At current pricing levels, a stripped 16-bit microprocessor (exclusive of random access memory) costs about \$800 in small quantities including some I/O control. Therefore the proposed arrangement will cost about \$3,200 in circuits with main memory yet to be added! From this result it seems far more advisable to build a 110-nsec DVT in TTL MSI which is known to be a far cheaper expedient and certainly an easier architecture to use.

4. In terms of IC count, each 16-bit elemental processor requires about 25 chips. Including main memory, the entire structure can be expected to require around 150-200 chips. However, many of the chips are 28- and 42-pin configurations. Hence, overall real estate savings are not improved as much as might be thought over a 300 can TTL realization of the standard DVT architecture.

A fully general structure, consisting of several parallel micro-processors, seems to be a losing proposition in terms of utility, cost, complexity, performance and form factor improvement. A better approach is to consider a somewhat specialized structure which retains a fair degree of flexibility through programmability. As an example, a processing structure based on the Markel LPC class of algorithms and employing 2 micro-processors is shown in Figure 10. The upper portion addresses the task of analysis. Straightforward real-time correlations are performed using special purpose digital hardware. But the less taxing (though conceptually more sophisticated) jobs of extracting filter parameters, coding/formatting information, and I/O supervision are programmed in a micro-processor. The pitch extraction path is also done partially in special purpose (analog) hardware and partially in the micro-processor. In the synthesis section, I/O supervision, decoding, buzz/hiss generation and vocal tract filter computations are done in the second micro-processor. A random access memory complex supplies code and working storage space to each micro-processor. Some of this, such as the encoding/decoding tables, could be common storage. The program memories should be independent, however.

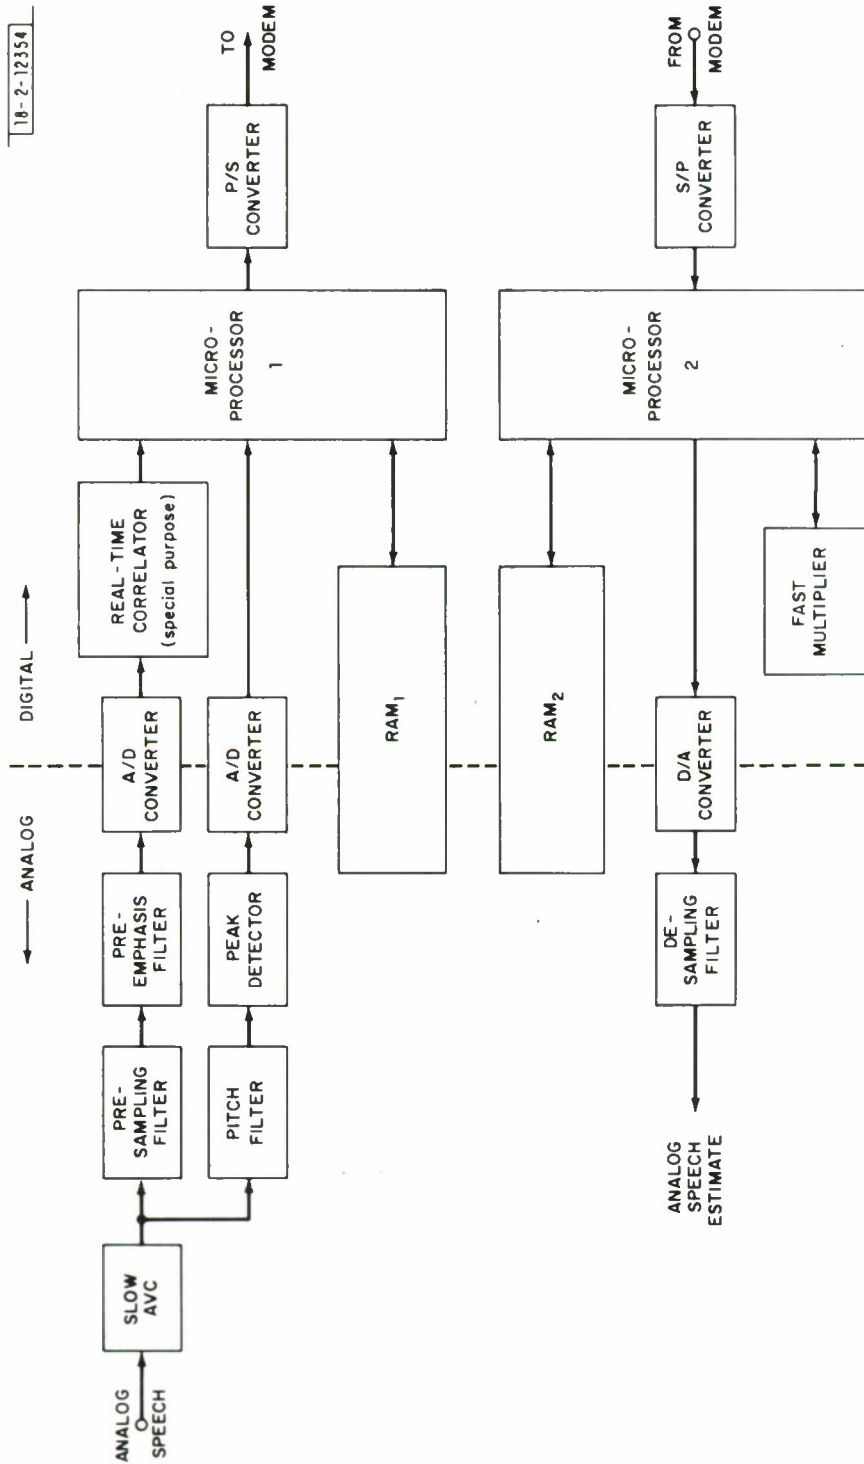


Fig.10. Semi-Programmable Speech Processor

Expected performance can be inferred from Table 1. Under synthesis it is seen that the DVT uses up about 13 of a 150- μ sec budget. A processor on the order of 10 times slower than the DVT doing only synthesis might use up 130- μ sec. It would seem that a comfortable margin relative to the 150- μ sec constraint is therefore maintained. In the analysis section the tasks of correlation and most of the real-time pitch analysis are done in external special purpose equipment. The remaining jobs need only be done once per frame implying that a processor 10 times slower than the DVT would have no real-time problems if confined to only these tasks. Thus it could perform other control tasks if desired.

The prospective IC count for such a structure does not seem unattractive either. Assuming a total RAM capacity of 2048 x 16, 2 16-bit micro-processors, and miscellaneous circuitry for the correlator and input/output traffic, a total count of well under 200 chips seems possible. The integrated circuit cost would be in the range \$2,500 to \$3,000. It must be realized that these figures are very tentative and very preliminary. Though promising, much more intensive, detailed studies of this class of micro-computer-based architecture must be conducted.

IV. Summary and Conclusions

In this report it was shown that the Lincoln Laboratory DVT design as it stands represents a very creditable set of tradeoff compromises when cost, size, performance, and utility are considered. The design was seen to be memory dominated in cost and IC count and as such could not be expected to benefit much from a custom LSI technology which did not address this issue. The irregularity of DVT structure implies definition of several unique LSI part types which, because of the high developmental cost per part type, serves to further discourage any more thoughts along this line. A hybrid packaging scheme is directly applicable to the memory problem, as well as to the rest of the miscellaneous logic comprising the machine, and at a much more tractable cost level. It is felt that this is the best route to cost and form factor improvement at the present DVT performance levels.

It was also seen that the new bipolar micro-processor chips by themselves yield results which are attractive in neither cost/performance, nor package count. For a fully programmable structure, a standard TTL MSI copy of the ECL DVT is a more effective approach. However, semi-programmable processor designs, addressing specific algorithm classes (such as LPC), may represent viable cost/performance alternatives with significant form factor improvement.

The four major design alternatives treated in the text are summarized in Table 7. The first 3 entries may be compared and contrasted as DVT-like structures starting with the current design and ending with a low performance, all-TTL copy of the ECL realization. It is also interesting to compare the last 2 entries, though not identical architectures, since they are both TTL systems. Two cost figures are given for each. The first represents an estimate of the recurrent Lincoln O.P. charges per unit (like Table 4). The second is an estimate of what similar costs might be for a

TABLE 7

SUMMARY OF DESIGN ALTERNATIVES

DESIGN	TECHNOLOGY	IC COUNT	SIZE (CU. FT.)	POWER (WATTS)	SPEED (NORM. TO REAL TIME)	RECURRENT LINCOLN O.P. COSTS	RECURRENT COMMERCIAL VENDOR COSTS
CURRENT DVT [*]	ECL MSI	476	1.25	225	2.0	\$13,100	\$8,200
	TTL MSI	213					
DOUBLE OVERLAP DVT (SLOWER MEMORIES)	ECL MSI	324	1.25	175	1.29	\$11,030	\$6,745
	TTL MSI	213					
CURRENT DVT IN TTL	TTL MSI	559	.8	180	~1.0	\$ 7,450	\$5,837
SEMI-PROGRAMMABLE USING TWIN INTEL μ -PROCESSORS	TTL MSI/LSI	<200	.3	<50	~1.0	\$ 6,550 ^{**}	\$3,000

* Including all peripherals.

** Assumes bipolar memory.

commercial vendor. It is seen that a commercial ECL DVT represents a very excellent buy if a flexible research tool is desired. However, for low performance, high production level applications, the μ -processor structure looks most attractive. Given the usual market pressures that come into play as new μ -processors become available, the cost projections can be expected to drop further. It would seem that the commercial market place is, for our purposes, the best mechanism for solving the cost problems of LSI yet reaping the obvious advantages.

REFERENCES

1. B. S. Atal and S. L. Hanauer, "Speech Analysis and Synthesis by Linear Prediction of the Speech Wave," J. Acoust. Soc. Am. 50, 637-655 (1971).
2. J. D. Markel, "Formant Trajectory Estimation from a Linear Least-Squares Inverse Filter Formulation," Speech Comm. Res. Lab., Santa Barbara, California, SCRL Monograph 7 (October 1971).
3. B. Gold and L.R. Rabiner, "Parallel Processing Techniques for Estimating Pitch Periods of Speech in the Time Domain," J. Acoust. Soc. Am. 46, 442-448 (1969).
4. B. Dunbridge, "LSI-Technology Overview," TRW Corporation, Redondo Beach, California, DCA Presentation, 10 July 1973.
5. H.S.E. Tsou, T.C. Berg and S.E. Hutchins, "High Speed LSI Processor for Voice Signal Processing," TRW Systems Group, Redondo Beach, California.
6. S.S. Husson, Microprocessing Principles and Practice (Prentice-Hall, New York, 1970)
7. J. Rattner, J-C. Cornet and M.E. Hoff, "Bipolar LSI Computing Elements Usher in New Era of Digital Design," Electronics Magazine (5 September 1974).

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ESD-TR-75-116	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Preliminary Investigation of Digital Speech Processor Hardware Implementations		5. TYPE OF REPORT & PERIOD COVERED Technical Note
		6. PERFORMING ORG. REPORT NUMBER Technical Note 1975-8
7. AUTHOR(s) Blankenship, Peter E.		8. CONTRACT OR GRANT NUMBER(s) F19628-73-C-0002
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lincoln Laboratory, M.I.T. P.O. Box 73 Lexington, MA 02173		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Program Element No. 30100
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Communications Agency 8th Street & So. Courthouse Road Arlington, VA 22204		12. REPORT DATE 5 February 1975
		13. NUMBER OF PAGES 40
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Electronic Systems Division Hanscom AFB Bedford, MA 01731		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES None		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) digital speech processor Digital Voice Terminal hybrid packaging microprocessor chip sets		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The current Lincoln Digital Voice Terminal design, analyzed with the intent of improving cost and form factor, is found to be memory dominant. Custom LSI and hybrid packaging are evaluated as to potential impact on the current design. Semi-programmable architectures based on commercially available bipolar microprocessors are viewed as near term solutions to the inexpensive mass-producible narrow band voice terminal design problem.		

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)